

CLAIMS

What is claimed is

- 1 1. A method of handling lock contention, the method comprising the steps of:
 2 a first requester transmitting to a lock management system a first request for a
 3 particular lock on a resource;
 4 wherein said lock management system manages locks on resources that may be
 5 granted to a plurality of processes that can access said resource;
 6 receiving a message from said lock management system that indicates that said first
 7 request to lock a resource is denied;
 8 wherein a blocking condition caused the denial of said first request;
 9 wherein said message includes first data; and
 10 based on said first data, said first requester transmitting a second request for
 11 notification that said blocking condition should no longer cause denial of a
 12 request for a lock on said resource.
- 1 2. The method of claim 1, wherein no process of said plurality of processes holds a lock
 2 issued by said lock management system for said resource.
- 1 3. The method of claim 1, wherein the steps further include:
 2 said first requester receiving said notification; and
 3 in response to receiving said notification, said first requester transmitting another
 4 request to said lock management system for said particular lock on said resource.
- 1 4. The method of claim 3, wherein the step of said first requester transmitting another
 2 request includes transmitting second data that indicates that said blocking condition
 3 should no longer cause denial of a request for said lock of said resource.

1 5. The method of claim 4, wherein the steps further include said lock management
2 system processing said other request without denial based on said second data.

1 6. The method of claim 1, wherein:
2 a second process of said plurality of processes is performing an operation that causes
3 said blocking condition;
4 said first data identifies another resource locked by said second process for which
5 said first requester may acquire a lock when said blocking condition should no
6 longer cause denial of a request for said lock of said resource; and
7 wherein the step of transmitting said second request includes transmitting a request
8 for said lock on said other resource.

1 7. The method of claim 1, wherein:
2 said first requester is a process of said plurality of processes;
3 wherein said resource is a data block in a b-tree index; and
4 wherein a second process of said plurality of processes is performing a block split
5 operation on said data block.

1 8. A method of managing locks by a distributed lock management system, the method
2 comprising the steps of:
3 a first lock manager on a first node receiving a first request for a first lock on a
4 resource from a first requester;
5 wherein said distributed lock management system includes said first lock manager;
6 determining that said first request may not be granted because of a blocking
7 condition;

said first lock manager storing in a data structure first data that may be used by said first requester to obtain notification that said blocking condition should no longer cause denial of a request for a lock on said resource; and said first lock manager transmitting to said first requester a first response that: indicates that said first request is denied, and includes a copy of said first data.

9. The method of claim 8, wherein the steps include: receiving a message that indicates that said blocking condition should no longer cause denial of a request for a lock on said resource; and modifying said data structure to indicate that said blocking condition should no longer cause denial of a request for a lock on said resource.

10. The method of claim 9, wherein the step of receiving said message includes receiving said message from said first requester.

11. The method of claim 8, wherein: the steps further include said first lock manager transmitting to another lock manager of said distributed lock management system a message requesting said first lock on said resource; and wherein the step of determining is based on a second response received from said other lock manager indicating that said first request cannot be granted; wherein said second response includes a copy of said first data.

12. The method of claim 8, wherein the steps further include: receiving a second request for another lock on said resource; determining, based on said first data, that said second request may not be granted;

4 said first lock manager transmitting to said second requester another response that:
 5 indicates that said second request is not granted, and
 6 includes a copy of said first data.

1 13. The method of claim 12, wherein:
 2 said first lock manager is a master of said resource; and
 3 wherein the step of receiving said second request includes receiving said second
 4 request from another lock manager.

1 14. The method of claim 12, wherein:
 2 said first lock manager and a process are on a node, wherein said process is different
 3 than said first requester; and
 4 the step of receiving said second request includes receiving said second request from
 5 said process.

1 15. The method of claim 8, wherein:
 2 said distributed lock management system includes a master for said resource; and
 3 wherein no lock is currently granted for said resource by said master.

1 16. A computer-readable medium carrying one or more sequences of instructions for
 2 handling lock contention, wherein execution of the one or more sequences of
 3 instructions by one or more processors causes the one or more processors to perform
 4 the steps of:
 5 a first requester transmitting to a lock management system a first request for a
 6 particular lock on a resource;
 7 wherein said lock management system manages locks on resources that may be
 8 granted to a plurality of processes that can access said resource;

9 receiving a message from said lock management system that indicates that said first
 10 request to lock a resource is denied;
 11 wherein a blocking condition caused the denial of said first request;
 12 wherein said message includes first data; and
 13 based on said first data, said first requester transmitting a second request for
 14 notification that said blocking condition should no longer cause denial of a
 15 request for a lock on said resource.

1 17. The computer-readable medium of claim 16, wherein no process of said plurality of
 2 processes holds a lock issued by said lock management system for said resource.

1 18. The computer-readable medium of claim 16, wherein the steps further include:
 2 said first requester receiving said notification; and
 3 in response to receiving said notification, said first requester transmitting another
 4 request to said lock management system for said particular lock on said resource.

1 19. The computer-readable medium of claim 18, wherein the step of said first requester
 2 transmitting another request includes transmitting second data that indicates that said
 3 blocking condition should no longer cause denial of a request for said lock of said
 4 resource.

1 20. The computer-readable medium of claim 19, wherein the steps further include said
 2 lock management system processing said other request without denial based on said
 3 second data.

1 21. The computer-readable medium of claim 16, wherein:

a second process of said plurality of processes is performing an operation that causes
 said blocking condition;
 said first data identifies another resource locked by said second process for which
 said first requester may acquire a lock when said blocking condition should no
 longer cause denial of a request for said lock of said resource; and
 wherein the step of transmitting said second request includes transmitting a request
 for said lock on said other resource.

22. The computer-readable medium of claim 16, wherein:
 said first requester is a process of said plurality of processes;
 wherein said resource is a data block in a b-tree index; and
 wherein a second process of said plurality of processes is performing a block split
 operation on said data block.

23. A computer-readable medium carrying one or more sequences of instructions for
 managing locks by a distributed lock management system, wherein execution of the
 one or more sequences of instructions by one or more processors causes the one or
 more processors to perform the steps of:
 a first lock manager on a first node receiving a first request for a first lock on a
 resource from a first requester;
 wherein said distributed lock management system includes said first lock manager;
 determining that said first request may not be granted because of a blocking
 condition;
 said first lock manager storing in a data structure first data that may be used by said
 first requester to obtain notification that said blocking condition should no
 longer cause denial of a request for a lock on said resource; and

13 said first lock manager transmitting to said first requester a first response that:

14 indicates that said first request is denied, and

15 includes a copy of said first data.

1 24. The computer-readable medium of claim 23, wherein the steps include:

2 receiving a message that indicates that said blocking condition should no longer cause

3 denial of a request for a lock on said resource; and

4 modifying said data structure to indicate that said blocking condition should no longer

5 cause denial of a request for a lock on said resource.

1 25. The computer-readable medium of claim 24, wherein the step of receiving said

2 message includes receiving said message from said first requester.

1 26. The computer-readable medium of claim 23, wherein:

2 the steps further include said first lock manager transmitting to another lock manager

3 of said distributed lock management system a message requesting said first

4 lock on said resource; and

5 wherein the step of determining is based on a second response received from said

6 other lock manager indicating that said first request cannot be granted;

7 wherein said second response includes a copy of said first data.

1 27. The computer-readable medium of claim 23, wherein the steps further include:

2 receiving a second request for another lock on said resource;

3 determining, based on said first data, that said second request may not be granted;

4 said first lock manager transmitting to said second requester another response that:

5 indicates that said second request is not granted, and

6 includes a copy of said first data.

1 28. The computer-readable medium of claim 27, wherein:
 2 said first lock manager is a master of said resource; and
 3 wherein the step of receiving said second request includes receiving said second
 4 request from another lock manager.

1 29. The computer-readable medium of claim 27, wherein:
 2 said first lock manager and a process are on a node, wherein said process is different
 3 than said first requester; and
 4 the step of receiving said second request includes receiving said second request from
 5 said process.

1 30. The computer-readable medium of claim 23, wherein:
 2 said distributed lock management system includes a master for said resource; and
 3 wherein no lock is currently granted for said resource by said master.

2025-10-01 10:00:00